

Multi-objective direct optimization of dynamic acceptance and lifetime for potential upgrades of the Advanced Photon Source

ANL/APS/LS-319

M. Borland, V. Sajaev, L. Emery, A. Xiao, ANL, Argonne, IL 60439, USA

(Dated: August 25, 2010)

Abstract

The Advanced Photon Source (APS) is a 7 GeV storage ring light source that has been in operation for well over a decade. In the near future, the ring may be upgraded, including changes to the lattice such as provision of several long straight sections (LSS). Because APS beamlines are nearly fully built out, we have limited freedom to place LSSs in a symmetric fashion. Arbitrarily-placed LSSs will drastically reduce the symmetry of the optics and would typically be considered unworkable. We apply a recently-developed multi-objective direct optimization technique that relies on particle tracking to compute the dynamic aperture and Touschek lifetime. We show that this technique is able to tune sextupole strengths and select the working point in such a way as to recover the dynamic and momentum acceptances. We also show the results of experimental tests of lattices developed using these techniques.

I. INTRODUCTION

As storage ring light sources push to smaller emittances and seek to accommodate the diverse requirements of users, the difficulty of achieving workable injection efficiency and Touschek lifetime continues to increase. Achieving workable injection efficiency requires sufficiently large dynamic acceptance (DA). Achieving workable Touschek lifetime requires sufficiently large local momentum acceptance (LMA).

The problem is clearest in the case of low emittance rings. Small emittance entails strong, frequent focusing elements in order to control the dispersion in the dipole magnets. This leads to higher natural chromaticity and smaller dispersion outside the dipoles, which entails stronger sextupoles. This increases the variations of tunes with amplitude, which makes achieving large DA difficult. Higher-order dependence of the tunes on momentum offset, combined with the need to run at positive chromaticity in order to increase the single-bunch stability threshold leads to a momentum-dependent tune footprint that may cross strong resonances, thus limiting the LMA and hence the lifetime. In modern light sources, designers incorporate a large number of independent sextupole families in order to provide the flexibility needed to tune both the DA and LMA.

We recently developed a direct (i.e., tracking-based) single-objective technique for tuning sextupoles in order to maximize DA and LMA [1, 2]. A single-objective variant of the method was successful in improving the operation of the APS. It was also used to explore potential symmetric upgrade lattices incorporating up to eight LSSs. Tests of a mock-up configuration created by turning off quadrupoles (as a proxy for removing them) showed excellent agreement with expectations.

In this paper, we discuss further development of the method along the lines discussed in previous publications, as well as application to nonsymmetric lattices. In addition to directly optimizing DA as before, we now directly optimize computed Touschek lifetime.

The remainder of this paper is organized as follows: First, we discuss desired features of the APS upgrade lattice, in particular, the provision of long straight sections. Next, we give background on storage ring characterization, e.g., dynamic acceptance, momentum acceptance, and so forth. We then explain the optimization method, followed by discussion of specifics of its application to a possible APS upgrade. Finally, we show the results of tests of mock-up lattices.

II. APS UPGRADE LATTICE ALTERNATIVES

The APS magnetic lattice consists of 40 double-bend sectors, normally operated in an essentially translation- and reflection-symmetric configuration. (The symmetry is slightly broken by the use of Decker distortion in the user sectors 1-35 [3], a feature that any upgrade must retain since it reduces backgrounds on x-ray beam position monitors.) Each sector has ten quadrupoles, seven sextupoles, and eight steering magnets per plane. In anticipation of the need to customize the electron optics, all quadrupoles and sextupoles have independent power supplies [4]. In addition to permitting optics correction [5], this has allowed provision of special reduced horizontal beamsize (RHB) optics for several beamlines. Its importance to the APS upgrade is two-fold: First, it permits us to entertain lattice customizations without the expense of additional power supplies and controls. Second, it permits us to mock-up many of these customizations before we commit to any hardware changes.

The most important of these customizations will be provision of long straight sections (LSSs) to (nominally) eight sectors. As described in [6], this will be accomplished by removing the long Q2 quadrupoles, one BPM, and one corrector on either side of the target straight section, then moving other components away from the straight section into the empty locations.

Given that there are 40 straight sections in total and that we will nominally have eight long straights, the most natural configuration is to make every fifth straight section a long one. This retains a high degree of symmetry in the lattice, which is traditionally considered necessary in order to obtain good dynamic and momentum acceptances. However, such a configuration is less than ideal. First, it requires relocation of beamlines to accommodate the locations of the LSSs, which would be expensive and would disrupt the beamlines' science programs. Second, because there are five consecutive straight sections used by the machine, one LSSs would be in some sense wasted. Finally, for implementation of the short-pulse x-ray (SPX) scheme based on crab cavities [7], we need to have two LSSs separated by a single short straight, since this permits accommodating the SPX cyrostats without excessively reducing the space available for undulators.

The last two of these concerns could be addressed with a configuration that features four groups of two LSSs separated by a standard straight. This could be done in a four-fold symmetric fashion. However, the expense of moving beamlines is still present and should be

avoided if possible.

Based on these considerations and recent progress in optimization of the nonlinear properties of lattices [1, 2], we have developed workable non-symmetric configurations incorporating eight LSSs.

III. BACKGROUND

A. Dynamic Acceptance

As alluded to above, in addition to designing appropriate linear optics, a storage ring designer must control the nonlinear dynamics of the lattice in order to obtain adequate dynamic and momentum acceptances. We'll begin with an explanation of the meaning and importance of these acceptances.

“Dynamic aperture” is characterized by the area in horizontal and vertical space into which particles may be injected and survive as stored beam, with the sole limitations of particle stability being those imposed by the particle dynamics. “Dynamic acceptance” is similar to dynamic aperture, but includes the effects of physical apertures. The dynamic acceptance is the most relevant quantity for our studies, since our physical apertures exist already and are well-defined. The importance of the dynamic acceptance is that it is a primary determinant of the injection efficiency. In the APS, we inject in the horizontal plane, i.e., with non-zero residual amplitude in the horizontal motion, so the horizontal dynamic acceptance needs to be much larger than the vertical. In addition, because we inject from inside the ring, the horizontal acceptance on the negative side is more important than on the positive side.

Dynamic acceptance (DA) is determined by tracking particles with increasing initial horizontal (x) and vertical (y) amplitudes until the boundary between survival and loss is found. This process is illustrated in Figure 1. A series of n line searches are conducted beginning at the origin. The coordinates along the lines are simply

$$x = R \cos \phi_n \tag{1}$$

$$y = R \sin \phi_n, \tag{2}$$

where $\phi_n = n\pi/L$ for an L -line search. For each line, we wish to find the smallest R , R_{da} , at which a particle will be lost within N turns. To do this, R is increased in small steps ΔR until

a loss is seen within N turns. (We search from small aperture to large in order to reduce the probability of being misled by a stable island at large amplitude, even though this requires more computational effort than searching from large amplitude to small.) Following this, we take one step back and the step interval ΔR is then subdivided to refine the determination of R_{da} . The set of (x, y) points resulting from the determination of R_{da} for each line defines the DA boundary.

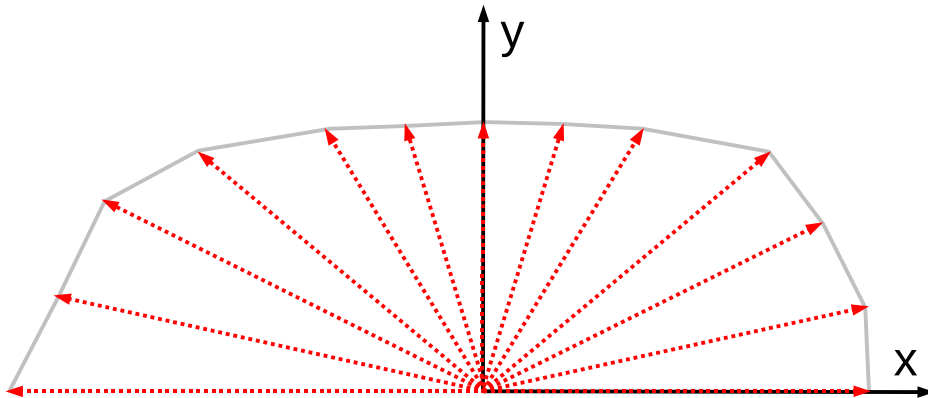


FIG. 1: Illustration of the method used to determine dynamic acceptance. Particles are launched at intervals along a series of lines beginning at the origin. After a loss is seen, interval subdivision is used to refine the determination of the minimum acceptance required for loss. The (x, y) location of first particle loss defines a point on the dynamic acceptance boundary.

B. Local Momentum Acceptance

Momentum acceptance is characterized by the maximum momentum displacement that a particle can undergo and still survive. In general, this quantity varies as a function of position around the ring, as well as for negative and positive momentum displacements. Hence, we are interested in the *local* momentum acceptance (LMA) characterized by two functions $\delta_n(s)$ and $\delta_p(s)$. These give, respectively, the maximum fractional negative and positive momentum displacements that a particle may experience at location s and still survive a chosen number of turns. These functions are primary determinants of the Touschek lifetime,

which is the dominant lifetime-limiting mechanism in the APS and other 3rd-generation light sources.

The method used to determine $\delta_{n,p}(s)$ [8] is similar to that used for determining the dynamic acceptance. At the exit of selected elements (i.e., at selected s positions), particles are given increasing positive or negative momentum kicks, then tracked N turns. When a loss is seen, we take one step back and proceed with a smaller step size to refine the determination of the boundary.

C. Particle Tracking Considerations

For both DA and LMA determination, particle tracking must include not only the effects of the magnets, but also longitudinal motion, radiation damping, physical apertures, and errors.

a. Longitudinal motion: The need for longitudinal motion (i.e., motion in the rf bucket) in the case of the LMA is clear, since the momentum acceptance may be determined either by transverse dynamics or the rf acceptance $\pm\epsilon_{rf}$. This is particularly the case with the APS, where we operate with relatively low rf voltage giving a $\epsilon_{rf} \approx \pm 2.3\%$. In the case of DA, we have observed some sensitivity to longitudinal motion, which results from the small path-length dependence on the betatron amplitude.

b. Radiation damping: Radiation damping must be included in the model for several reasons. First, it may increase the DA or LMA by damping the growth of particle motion near the stability boundary. Second, it stabilizes the determination of the acceptance by sweeping the particle's amplitude over a small range. This ensures that an unstable amplitude is not “stepped over” in the search. For this reason, in computing the LMA it is advisable to choose a momentum step size $\Delta\delta$ that is smaller than $\epsilon_{rf}(1 - e^{-NT_r/\tau_\delta})$, where $T_r = 3.68\mu\text{s}$ is the revolution time and $\tau_\delta = 4.8\text{ms}$ is the longitudinal damping time. For APS, we typically take $N = 400$, which means that over 25% of the initial momentum displacement decays away in the course of the simulation. A third reason for including damping in LMA simulations is that this large momentum slew potentially sweeps the particle past many resonances, the effect of which would be missed if radiation damping was omitted. Similar reasoning applies for DA determination, where radiation damping results in a rapid drop in amplitude that may result in sampling resonances due to the tune shift

with amplitude.

c. Physical apertures: Although we commonly speak of the “dynamic aperture,” what we are in fact interested in is the smaller of the dynamic and physical acceptances. Often, we find that the effective aperture is smaller than the physical aperture, which may lead to the erroneous conclusion that the physical aperture is not important. Including small vertical physical apertures is particularly important when there is coupling of horizontal and vertical motion, as it can result in significant reductions in the horizontal acceptance. All simulations include realistic physical apertures unless otherwise stated.

d. Inclusion of errors: Inclusion of errors refers not only to multipole errors resulting from magnet imperfections, but, more importantly in the case of APS, to random errors in the strength and alignment of magnets. For example, without errors the effect even of the half-integer resonance may not be seen in a tracking simulation. Similarly, without errors we would not have coupling between the horizontal and vertical planes, which is well known to strongly affect the dynamic acceptance of light sources with small vertical gap insertion devices. Hence, we must include lattice errors in the simulations.

Effective methods exist for correcting linear optics [9] and coupling (e.g., [10–12]), which is important in light sources because of the small insertion device vertical apertures. In the APS, for example, we correct lattice function errors to the 1% rms level [5] and coupling to the 1% level. These are commonly achieved in modern storage rings.

It would seem that in order to be realistic, we must not only include errors, but also correct those errors using real-world techniques. However, correction is not essential in the simulations. Once correction is performed, we are left with residual errors of a certain level. Instead of modeling relatively large errors with correction, we can simply use smaller random errors that give lattice function and coupling errors at post-correction levels. This neatly side-steps a considerable complication. We’ve found that assuming quadrupole and sextupole strength errors at the 0.02% rms level and tilt errors at the 0.05 mrad rms level gives lattice errors at or above the desired level. It is then only necessary to apply a coupling correction algorithm to emulate our operational method of setting of the vertical emittance.

IV. OPTIMIZATION TECHNIQUE

The properties described in the last section depend in detail on the linear and nonlinear tuning of the lattice. A primary determinant of these properties is the configuration of sextupole strengths. A very common approach [13] to optimization of sextupole settings is to minimize various resonance driving terms, which in turn minimizes tune shifts with transverse amplitude $d^n\nu_{x,y}/dA_{x,y}^n$ and momentum offset $d^n\nu_{x,y}/d\delta^n$. However, as described below, we operate the APS with chromaticities $d\nu_{x,y}/d\delta$ that are considerably different from zero. Hence, the utility of minimizing the higher-order chromaticities is doubtful. Indeed, we very likely need to employ significant higher-order terms to minimize the size of the tune footprint over the desired momentum acceptance.

The issue is less obvious in the case of the amplitude-dependent tune shifts. However, experience has shown that small amplitude-dependent tune shifts are not always optimal, depending on the proximity of the working point to strong resonances. As is known, a large tune shift with amplitude can defuse a resonance by preventing quasi-stable particles from remaining on the resonance.

Even when it is applicable, the traditional method of sextupole optimization does not guarantee good DA and LMA. If it did, we would not need tracking simulations to validate the sextupole configuration, and yet such simulations are universally used. Informed by these observations, we have developed a new technique [1, 2] for sextupole optimization that relies on tracking simulations directly. It is made possible by the availability of relatively inexpensive computing clusters and fully-scriptable accelerator simulations.

Although our method could use any tracking code, the ability to create fully scripted simulations is essential, since matching and tracking must run without human intervention. Thus, we use the tracking program **elegant** [14, 15], as well as the SDDS Toolkit [20] and **geneticOptimizer** [16].

In this method we use many computers simultaneously to evaluate the DA and LMA for various lattice tunings (e.g., tunes and sextupole settings). As described above, DA and LMA computation includes radiation damping, synchrotron oscillations, and physical apertures. After completion of a sufficient number of evaluations, a genetic algorithm is used to “breed” more candidate configurations based on the best configurations seen so far. The process continues until a sufficiently good solution is obtained or until the results stop

improving.

As mentioned above, it is important to include errors in the DA and LMA simulations. In order to have a hope that the optimization will converge, it is of course essential that the errors remain unchanged for each evaluation. Hence, we optimize with a single, fixed ensemble of errors. It is always possible that we may choose an error ensemble that happens to provide atypically good results. We guard against this possibility by evaluating the optimized configuration with a large number of ensembles.

Any optimization requires one or more penalty functions. In the work reported here, we used a multi-objective optimization. The objectives are derived from the DA and lifetime.

A. Dynamic Acceptance Penalty Function

While the DA is frequently computed, quantitative use of the acceptance is uncommon and can be fraught with uncertainty. Typically, we quote the horizontal acceptance as a single number, but do not quantify the shape of the acceptance. Such a number could well be highly misleading. However, in order to optimize DA, we must obtain one (or a few) quantities that clearly indicate how good a solution is. Our approach is to use the area of the DA, but to compute it with certain restrictions that ensure it is a useful quantity.

One potential pitfall of using the DA area is that the optimizer might provide a useless solution with large vertical acceptance and small horizontal acceptance. This can be prevented by simply not scanning the vertical coordinate beyond the minimum requirement. This prevents the algorithm from “seeing” any vertical DA that exceeds requirements.

Another potential pitfall is that the acceptance may not have a smooth shape. In some instances, depending for example on the strength of coupling resonances, one might find useless protrusions that enlarge the area but that are not useful for injection. We avoid this by applying a clipping algorithm to the DA boundary, as illustrated in Figure 2.

A third potential pitfall is that the acceptance may be enlarged at $x > 0$ at the expense of $x < 0$ (at APS, we inject on the $x < 0$ side). We avoid this by clipping off any DA that extends beyond $x = 7$ mm, which provides generous room for oscillation of the stored beam due to our mismatched kicker bump.

With these restrictions and modifications, we’ve found that the DA area A_d is a robust indicator of a good solution. Because we wish to maximize this area, the related objective

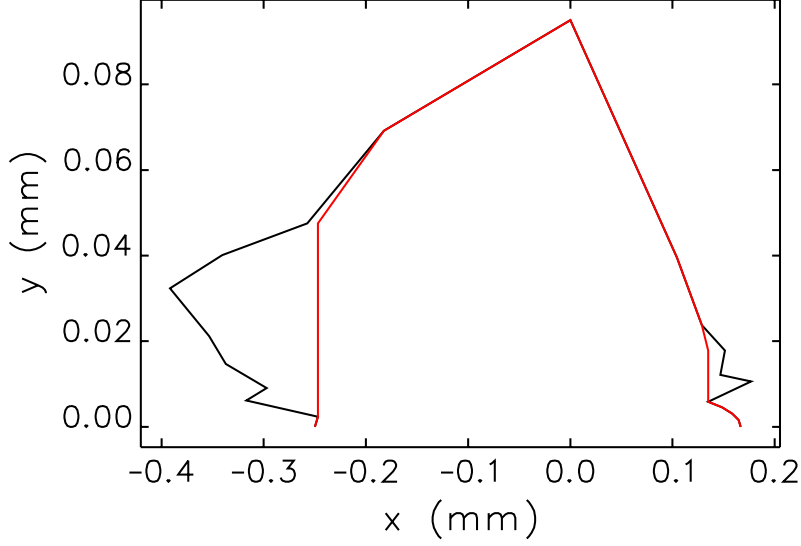


FIG. 2: Illustration of the behavior of dynamic acceptance boundary clipping method. The black line shows the original dynamic acceptance boundary, while the red line shows the clipped boundary.

function is $-A_d$.

B. Lifetime Penalty Function

The Touschek lifetime can be computed from the LMA, the Twiss parameters, the beam emittances, and the bunch charge. We do this using the program `touschekLifetime` [18]. This program uses the Twiss parameters and LMA computed by `elegant` and applies Piwinski's formalism[19].

In order to reduce the computer resource requirements, we do not compute the LMA at the exit of every element. Instead, we compute the LMA at the exit of the S1, S3, and S4 sextupoles in the first six (of 40) sectors. Experience shows this is representative of the LMA around the ring. In order to compute the Touschek lifetime, we simply assume that the LMA repeats the same pattern in the subsequent 34 sectors. We check out results in ensemble evaluation by determining the LMA at the exit of all sextupoles.

The objective function related to the Touschek lifetime τ_T is $-\tau_T$, since we want to maximize the lifetime.

In many cases, optimization of the lifetime results in a relatively small momentum-dependent tune footprint that avoids major resonances, such as the integer, half-integer,

and main coupling resonances. However, in some cases we have seen stable half-integer resonance crossings, which apparently are possible because of strong coupling between the planes and/or large tune shift with amplitude. Observation of phase space trajectories for such crossings shows an exchange of amplitude between the planes, as well as mitigation of the effects of the resonances due to fairly rapid synchrotron motion and, presumably, large tune shifts with momentum and amplitude.

In spite of this, we elected to disallow crossing the half integer resonance. We did this by modifying the LMA computation to “cap” the momentum acceptance at the momentum deviation that corresponds to the resonance crossing, if any. This requires that we first track a series of off-momentum particles to determine their tunes. This adds only a small amount of additional running time. The momentum acceptance cap is accepted as an additional parameter by `touschekLifetime`.

C. Tune Variation

Variation of the tunes is important in finding an optimal solution. In the case of complex lattices with many different types of sectors, this can be challenging. In the present APS lattice, we have three types of sectors: ordinary user sectors, which include Decker distortion [3]; special user sectors with reduced horizontal beamspace (RHB); and non-user sectors, which do not have Decker distortion. Of the 40 sectors in the ring, a least 33 are ordinary user sectors while five are non-user sectors. Changing the working point is thus relatively simple, even using a quadrupole knob for the ordinary sectors.

When the lattice becomes more complex, this is not true. In particular, with long straights we have reduced ability to change the tunes while satisfying other constraints, such as maximum beta functions, lattice functions at straight sections, and uniformity of phase advance, not to mention preventing the effective emittance from increasing appreciably. For complex lattices, the most robust approach is to perform matching of the entire ring at once, which can be time-consuming. If we must perform such matching for each configuration during the optimization, it can easily add 25% to the running time.

Instead, we chose to make a grid of solutions covering the tune plane on $\nu_x : [36.1, 36.4]$ and $\nu_y : [19.1, 19.4]$ with a spacing of 0.01. We started by creating a single solution by matching individual sector types “by hand.” We then gradually covered the tune plane

by working our way outward from this starting solution, taking steps of no more than $\nu_x = 0.01$ and $\nu_y = 0.01$. This can be done on a cluster in a matter of a few hours, depending on the difficulty of the lattice. Once we have this grid of solutions, we create a table of quadrupole strenghts as a function of tunes. We can then create new solutions by performing two-dimensional interpolation as a function of the tunes. Tests showed that the resulting interpolated solutions are very well behaved. This process was made relatively straightforward by the fact that `elegant` saves and loads lattice data in self-describing data sets (SDDS) files [20–22], which allowed collation and interpolation of the quadrupole strength table without writing any new code, using the SDDS toolkit programs.

D. Details of the Algorithm

Having reliable objective functions is perhaps the most important part of any optimization. This has been covered in the previous subsections. Here, we discuss several details of the optimization algorithm itself, which is implemented in the `geneticOptimizer` Tcl script [16]. The algorithm is as follows:

1. Create N randomized configurations. Typically, these are “small” perturbations from a reasonable starting point, e.g., a previous similar configuration or the present operational configuration. N varies depending on the computing resources available. Typically, we have $N \sim 100$, but have used $N > 1000$ for some of the work reported here.
2. Submit N jobs to a cluster to evaluate the configurations.
3. Wait until at least M (e.g., 4 to 10) configurations are newly-completed. (Note that the running time is not the same for all configurations, since a small DA/LMA takes less time to simulate.)
4. Collect information on all completed configurations.
5. If the best configuration is adequate, stop. Otherwise
 - (a) Perform a non-dominated sort to find the first-rank solutions [17].
 - (b) Determine how many unused slots P are available for new computing jobs.

- (c) For each open slot:
 - i. Randomly select two configurations from this group to act as “parents.”
 - ii. Randomly blend the attributes of the parents to “breed” a new configuration.
 - iii. Submit a new job to evaluate this configuration.
6. Wait for a least one additional job to complete.
7. Return to step 4.

The quantities N and M have been determined in only a rough fashion based on experience, rather than any systematic program of study.

V. APPLICATION TO APS UPGRADE

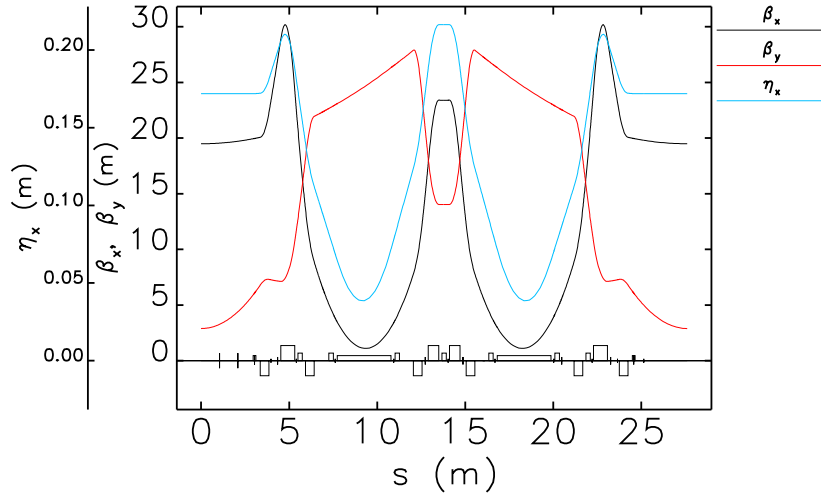


FIG. 3: Present lattice functions for the first sector of the APS.

The present APS lattice functions are shown in Figure 3. Any new lattice design must satisfy a set of constraints, which are of course well-satisfied by the present lattice. These constraints include

- Maximum beta functions in the horizontal and vertical planes of 40 m and 30 m, respectively. Ideally, maximum beta functions should be as small as possible, particularly at ID transitions, in order to reduce the effective chamber impedance.

- Horizontal beta functions at the center of ID straights of between 15 and 25 m. In general, values close to 20 m are used now and are preferred.
- Vertical beta functions at the center of the ID straights within 20% of $\beta_{\text{opt}} = L/2$, where L is the length of the small gap chamber. This maximizes the vertical acceptance. The beta function β_e at the end of the chamber is relatively insensitive to the exact value $\beta_c = \beta_{\text{opt}} + \Delta\beta_c$ at the center, being governed by

$$\frac{\Delta\beta_e}{\beta_e} \approx \frac{1}{2} \left(\frac{\Delta\beta_c}{\beta_c} \right)^2 \quad (3)$$

- At present the APS effective emittance in the straight sections is 3.1 nm, resulting from a 2.5 nm natural emittance combined with the presence of dispersion in the straight sections. Because the effective emittance has been optimized to the present low value, any changes are likely to result in an increase. This increase should be kept to under $\sim 10\%$. In order to keep the emittance close to its present value, the horizontal integer tune should be 36 (the present value) or higher. There is little reason to raise it above this value, so we've elected to keep both tunes close to the present values, $\nu_x = 36.15$ and $\nu_y = 19.25$.
- Because of the need to run with high single-bunch current and the present lack of a bunch-by-bunch feedback system, APS operates with high chromaticities, which reduces the impact of the transverse mode coupling instability. In 24-bunch mode, we typically have $\xi_x \approx 7$ and $\xi_y \approx 6$, while in hybrid mode we require $\xi_x \approx \xi_y \approx 11$. We anticipate that an existing prototype single-bunch feedback system will transition into routine operations in the near future. With this system, the requirement for exceeding 16 mA per bunch should be reduced to $\xi_x \approx \xi_y \approx 8$ [23], while $\xi_x \approx \xi_y \approx 7$ will be needed for 24 bunch mode at 200 mA. Hence, we have targeted $\xi = 8$ in our designs. (Note, however, that this gives only a small margin between the accumulation limit and the present 16 mA hybrid mode operating current. Ideally, chromaticity should be raised to 9 in both planes. This may be mitigated by further improvements to the feedback system.)
- We do not intend to replace any main accelerator magnets or power supplies as part of the upgrade. Hence, we must respect existing power supply limits, which give

$K_1 \leq 0.9 \text{ m}^{-2}$ for quadrupoles. For sextupoles, we are limited to $K_2 \leq 21.6 \text{ m}^{-3}$ for most sextupoles. We can if necessary add inexpensive “noses” to the sextupole poles as was done for the S2 family, raising the limit to 30.4 m^{-3} .

- As mentioned above, sufficient dynamic acceptance is necessary in order to have efficient injection. Simulations of the present operating lattice indicate that our dynamic acceptance is -12 to -15 mm on the injection side, as shown in Figure 4. Hence, we attempt to achieve this value in optimization of new lattices. The dynamic acceptance on the positive side is less important; we attempt to achieve at least 7 mm. In the vertical plane, we target 1 mm, which is quite generous compared to the typical rms beam size of $10 \text{ } \mu\text{m}$ for the stored beam and 0.2 mm for the injected beam (assuming 10% coupling in the booster).

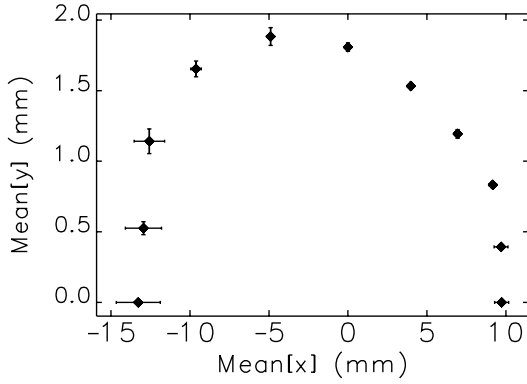


FIG. 4: Nominal dynamic acceptance computed for 50 error ensembles for $\xi_x = 7$ and $\xi_y = 6$. The error bars show the standard deviation of the DA boundary location.

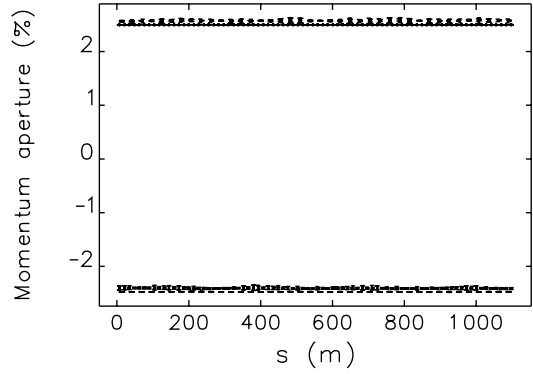


FIG. 5: Nominal momentum acceptance computed for 50 error ensembles for $\xi_x = 7$ and $\xi_y = 6$. The error bars show the standard deviation of the boundary location. An rf voltage of 9 MV is assumed, and is the limiting factor in determining the momentum acceptance.

- Although APS runs in top-up mode [24] and can therefore tolerate somewhat short beam lifetime, there are limits both in the amount of charge and frequency of injections. Ideally, we would inject no more frequently than every 120 seconds, although in hybrid mode at present we inject every $\Delta T_i = 60 \text{ s}$. Although, our injector is nominally

capable of delivering $\Delta Q_i = 6$ nC per shot, a value of 4 nC is more reliable. The minimum workable lifetime is

$$\tau_{\min} = \frac{Q_s \Delta T_i}{\Delta Q_i \eta_i}, \quad (4)$$

where Q_s is the total stored charge and η_i is the injection efficiency. At 200 mA and 80% efficiency, this gives $\tau_{\min} = 3.8$ h for a 60-second top-up interval.

Another consideration is the droop in the charge in each bunch between injections, which should be kept to a minimum. This is given by

$$D = 1 - e^{\frac{-B \Delta T_i}{\tau}}, \quad (5)$$

where B is the number of stored bunches. If the lifetime is only 3.8 hours for 24 bunches, we would have $D = 0.1$, or a 10% droop. This is probably acceptable, but clearly having a longer lifetime is desirable, if for no other reason than to prevent radiation damage of components. Hence, we target a minimum lifetime of 5 hours.

The lifetime is strongly affected by the LMA. The reference LMA for the present APS lattice is shown in Figure 5. This is sufficient to provide lifetimes of 8 to 9 hours at 100 mA in 24 bunches for our nominal coupling of 1.3%.

We began with optimization of a configuration with eight LSSs in locations thought to be convenient for planned beam line upgrades. In particular, LSSs were placed in straight sections 1, 6, 8, 11, 13, 16, 28, and 33. Since the LSSs are in essentially random, rather than symmetric, locations, we refer to this lattice as 8RLSS. The linear optics is reflection symmetric for the two sectors on either side of each LSS. For example, sector 1 (a short-to-long sector) and sector 2 (a long-to-short sector) quadrupoles are mirror images of each other. As noted above, there are seven sextupoles per sector. Although these are in symmetric physical locations in the sector, the strengths are allowed to vary without this symmetry constraint, which has been found to be essential in getting good solutions [1]. The sextupoles are, however, the same in all sectors with the same sequence of quadrupole strengths (as seen by the beam). Hence, the sextupoles in sectors 1, 6, 8, 11, and so on are identical, as are those in sectors 2, 7, 9, 12, and so on. Since there are seven sextupoles per sector, we have a total of 14 independent sextupole variables for the long-to-short and short-to-long sectors.

In addition to the sectors surrounding the LSSs, we need solutions for ordinary sectors having short straights at each end. Two such solutions are needed, one for the (normal) Decker distorted sectors and one for the non-Decker sectors (which are used for machine operation and have no beamlines). Although the linear optics in these sectors are very similar, we allowed the sextupoles to vary independently. This provides 14 additional sextupole variables. Two of these variables are used to set the chromaticities. In total, then, we have 26 free sextupole variables available for nonlinear dynamics optimization. As mentioned above, we also allow the tunes to vary, bringing the total number of variables for this optimization to 28.

Starting from the nominal, symmetric sextupoles for a symmetric linear solution, we optimized these sextupoles for target chromaticities of $\xi_x = \xi_y = 3$. Once a satisfactory solution was obtained, we raised the chromaticity in steps $\Delta\xi = 2$ up to $\xi = 7$. Following that, the optimization becomes more challenging and smaller steps are needed. At present, solutions exist up to $\xi_x = \xi_y = 8$. Figures 6 and 7 show the linear optics for this configuration, while Figures 8 and 9 show the dynamic and momentum acceptances.

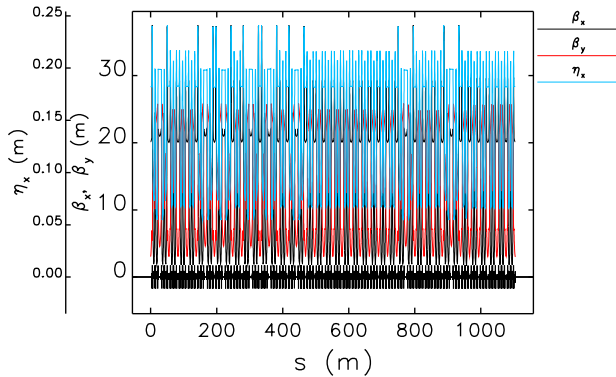


FIG. 6: Optical functions for the full ring with eight LSS at “random” locations.

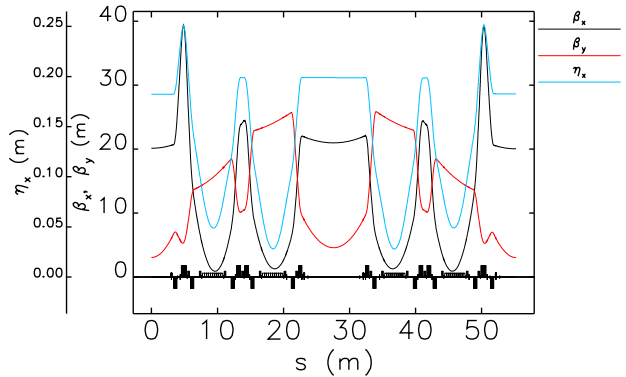


FIG. 7: Detail of optical functions for solution with eight LSS. Here we show the functions for the first two sectors, which starts and ends with short straights while having a long straight in the center straight.

It is of interest that the dynamic acceptance in this lattice is enlarged on the negative x side compared to the lattice without long straight sections. Indeed, the dynamic acceptance

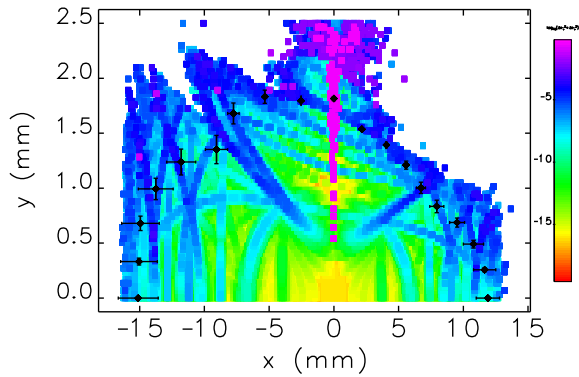


FIG. 8: Dynamic acceptance from 50 ensembles for the lattice with eight LSS and $\xi_x = \xi_y = 8$, superimposed on the frequency map for the machine with no errors. The frequency map is color-coded by the diffusion, defined as $\log(\Delta\nu_x^2 + \Delta\nu_y^2)$.

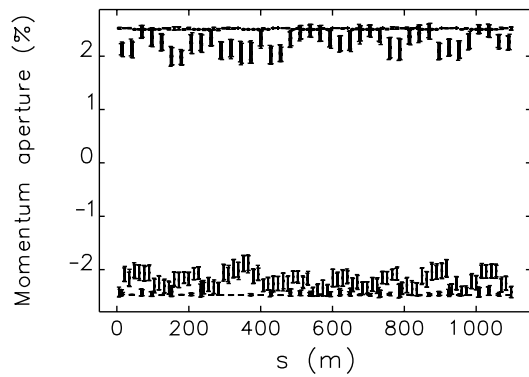


FIG. 9: Momentum acceptance from 50 ensembles for the lattice with eight LSS and $\xi_x = \xi_y = 8$.

is larger than the nominal limiting physical acceptance, which results from a curious distortion of the phase space ellipses between the injection point at the end of sector 40 and the location of the minimum aperture (which is -15 mm in the horizontal at straight section 4). This could only be brought about through adjustment of a significant number of sextupoles between the injection and limiting aperture points, which is what we have in the present case as Figure 10 shows.

Since a workable 8RLSS lattice has been obtained, we move next to add the required optics for the short-pulse x-ray scheme, known as SPX. The center of the SPX system will be straight section 7. Because of the length of the cryostats for the deflecting cavities, we need long straights in sectors 6 and 8. This not only allows accommodating the cryostats but still permits sufficient space for insertion devices in those sectors. The vertical betatron phase advance between the crab cavities in straights 6 and 8 must be 360° . Fortunately, the nominal vertical tune for the APS is 19.3, so we are close to this value already.

There are four possible placement options for the cavities. For example, we could have cavities in the upstream half of straights 6 and 8, which we refer to as US+US. We could also have DS+DS, DS+US, and US+DS. It is likely that the desired configuration will evolve

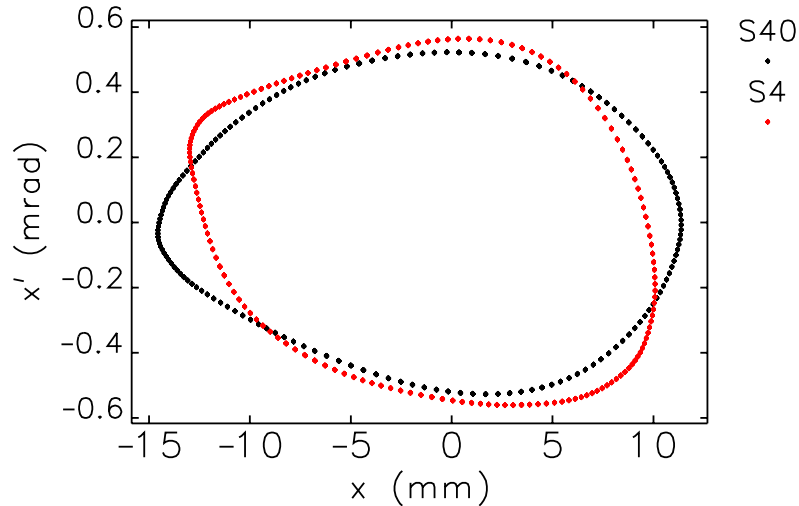


FIG. 10: Distortion of horizontal phase space in the eight LSS lattice with $\xi_x = \xi_y = 8$. Shown are the particle coordinates at the ends of sectors 40 (injection point) and 4 (minimum aperture point). The phase space is distorted in such a way that the effect of the limiting horizontal aperture of -15 mm at S4 is diminished.

in time as more or fewer experiments demand access to chirped pulses. We present results here for the US+US configuration, which involves the least most significant changes in the optics. Work is underway on the more challenging DS+US configuration.

We started with the solution from the 8RLSS optimization and determined linear optics to satisfy the SPX requirements with tunes that are close to the previous optimum values. We then performed optimization of the 14 interior sextupoles (i.e., those in sectors 7 and 8) to control vertical emittance growth due to the crab cavities [25]. We also developed linear optics for a tune grid, as described above. Finally, we performed optimization of the exterior sextupoles. The interior sextupoles are fixed, which was found to be acceptable given the small changes in the linear optics of sectors 7 and 8 that are introduced as we change the overall tunes. To obtain good results, it was necessary to allow the 14 sextupoles in sectors 6 and 9 to vary, giving a total of 40 independent variables. (This was reduced from 42 by tying together the chromatic sextupoles in the Decker and on-Decker sectors, which seemed to aid convergence.)

The linear optics for the region surrounding the SPX straight is shown in Figure 11. Figures 12 and 13 show the dynamic and momentum acceptance for 50 ensembles with

$\xi_x = \xi_y = 7$. One interesting feature is that the dynamic acceptance is significantly larger than before, a result of the large number of sextupoles available for adjustment. However, the momentum acceptance is not as good, which is why the solution is presently limited to $\xi = 7$.

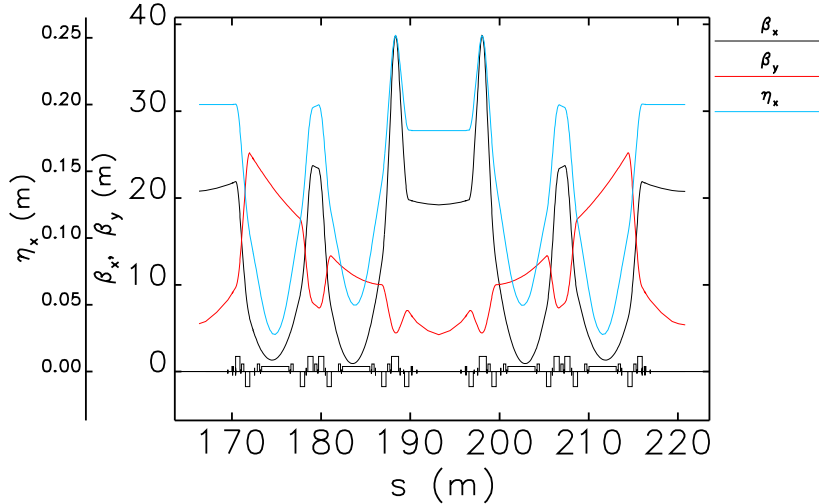


FIG. 11: Detail of optical functions for solution with eight LSS and SPX. Here we show the functions for the first two sectors, which starts and ends with short straights while having a long straight in the center straight.

Finally, we have performed optimization for the 8RLSS + SPX lattice with the addition of reduced horizontal beam size (RHB) in one location (sector 20). Using the quadrupoles in sectors 20 and 21, we adjusted the linear optics to target a horizontal beamsize of $150 \mu\text{m}$, which is about 50% of what is delivered in a normal straight section. The sextupoles in sectors 20 and 21 were allowed to vary independently, giving a total of 49 independent variables for the DA/LMA optimization. Figures 14 and 15 show the dynamic and momentum acceptances, respectively, for chromaticities of 7. While the dynamic acceptance is acceptable, the momentum is noticeably reduced. Again, this has so far prevented raising the chromaticities to the desired value of 8.

Table I compares the parameters of the four lattices. We see that the effective emittance has increased by less than 10% in going from the present lattice to the final lattice. The split between the horizontal and vertical tunes is generally larger than is typical in operations.

Table II gives a comparison of the lifetime for the four lattices for various fill patterns.

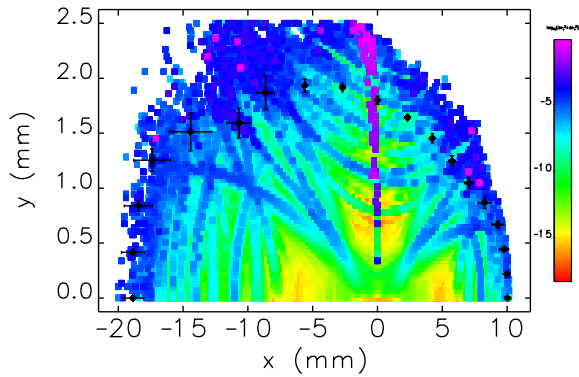


FIG. 12: Dynamic acceptance from 50 ensembles for the lattice with eight LSS, SPX, and $\xi_x = \xi_y = 7$, superimposed on the frequency map for the machine with no errors. The frequency map is color-coded by the diffusion, defined as $\log(\Delta\nu_x^2 + \Delta\nu_y^2)$.

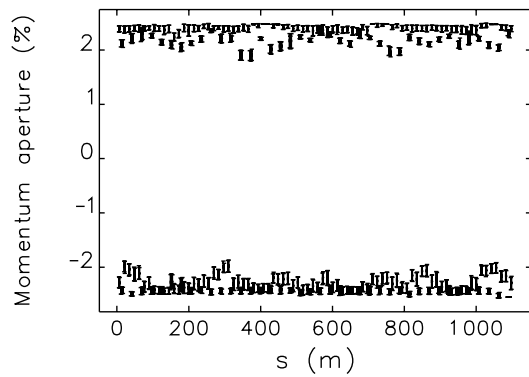


FIG. 13: Momentum acceptance from 50 ensembles for the lattice with eight LSS, SPX, and $\xi_x = \xi_y = 7$.

For each lattice, the lifetime is computed for each of the 50 ensembles. Variations in lifetime result from variations in the momentum acceptance and the effectiveness of coupling correction. The median gives a prediction of the most likely lifetime in operations, while the 5th percentile gives a prediction of the lowest likely lifetime. That is, the probability is 95% that we will have lifetime greater than this value. With the exception of the lattice incorporating RHB, all the lattices have median lifetimes above our goal of 5 hours for the 24 bunch patterns. The lattice with RHB is marginal at 200 mA, but may prove workable with further optimization.

We list the 200 mA, 324 bunch mode evaluated with the high chromaticity required for the 24 bunch mode. The lifetime is quite long in spite of the excessively high chromaticity, which provides an operational convenience as it obviates the need to change sextupoles when changing between the 24 and 324 bunch fill patterns. Presently, 324 bunch mode does not require top up. However, if we reduced the vertical emittance to, say, 10 pm, the lifetime would drop to about 20 hours, which would require occasional top-up. This mode would provide significantly higher brightness than the other modes, owing to the low coupling.

TABLE I: Comparison of major parameters of the four lattices at 7 GeV

	Now	8RLSS	8RLSS+SPX	8RLSS+SPX+RHB	Units
	$\xi_{x,y} = 7, 6$	$\xi = 8$	$\xi = 7$	$\xi = 7$	
ν_x	36.127	36.117	36.102	36.100	
ν_y	19.248	19.315	19.250	19.293	
Maximum β_x	29.613	36.316	36.297	36.308	m
Maximum β_y	25.705	25.803	25.481	30.055	m
Maximum η_x	0.224	0.233	0.243	0.261	m
Average β_x	13.656	13.810	13.768	13.641	m
Average β_y	15.173	14.491	14.076	13.890	m
Average η_x	0.156	0.152	0.157	0.154	m
Standard ID Straight Sections					
Effective emittance	3.186	3.305	3.344	3.391	nm
β_x	20.092	20.148	20.091	20.092	m
η_x	0.189	0.176	0.184	0.179	m
β_y	3.088	3.059	3.146	3.222	m

VI. EXPERIMENTAL TESTS

Because the APS has independent power supplies for all quadrupoles and sextupoles, it is possible to mock up a long straight section. This is done by setting to zero the power supplies for the Q1 magnets on either side of the target straight section. This reduces the quadrupole strength to $K_1 = -0.006 \text{ m}^{-2}$, compared to a maximum strength of 0.9 m^{-2} . Using the same techniques as discussed in the previous section, we can optimize the linear and nonlinear properties of the lattice to obtain a configuration with large predicted dynamic acceptance and lifetime. Such optimizations were carried out for a series of mock-up lattices corresponding to each of the configurations described in the previous subsection. In addition, we have experimentally tested two of the mockup lattices: 8RLSS and 8RLSS+SPX+RHB.

Here we need to mention that it has been APS policy from the beginning of operation to allow large user-requested beam steering. Over time the storage ring and beamlines

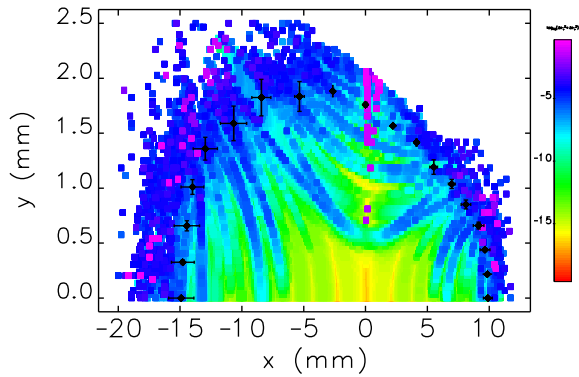


FIG. 14: Dynamic acceptance from 50 ensembles for the lattice with eight LSS + SPX + RHB and $\xi_x = \xi_y = 7$, superimposed on the frequency map for the machine with no errors. The frequency map is color-coded by the diffusion, defined as $\log(\Delta\nu_x^2 + \Delta\nu_y^2)$.

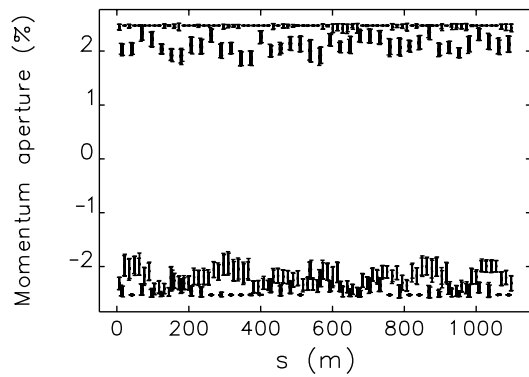


FIG. 15: Momentum acceptance from 50 ensembles for the lattice with eight LSS + SPX + RHB and $\xi_x = \xi_y = 7$.

have settled, and at many locations the steering has accumulated to several millimeters of orbit distortion. This orbit is called “user orbit” and deviates from the centers of magnets significantly. This would have been a big problem for the storage ring optics, but fortunately the APS has separate power supplies for all quadrupoles (and sextupoles). As long as the optics is regularly corrected there seem to be no significant negative effects from operating on a significantly non-zero orbit.

The first step in our experiments was to test the simpler 8RLSS lattice. We started by setting this lattice up on the “user orbit.” Lattice setup includes installing new quadrupole and sextupole setpoints, correcting the orbit, measuring and correcting the beta functions and dispersion, adjusting the coupling, and optimizing the injection efficiency. In this experiment, we found that we could not improve injection efficiency above 40% or store beam current above 40 mA. We performed a lifetime vs betatron tunes scan and found a strong resonance line at $2\nu_x - \nu_y = 54$. This is a skew-sextupole-type resonance that is excited when large coupling is present.

To reduce this resonance, we next set the lattice up on the “reference orbit,” i.e. the orbit going through the centers of the magnets. We repeated the lifetime vs. tunes scan

TABLE II: Comparison of computed lifetimes for various lattices and fill patterns. For each fill pattern, the median and 5th percentile lifetimes are shown for all lattices. The median vertical emittance is 35 pm. Note that $\xi = 7$ is insufficient for 16 mA single bunch. In addition, 324-bunch mode does not require significant chromaticity, but operationally it is convenient to use the same values as used for 24-bunch mode.

	Now	8RLSS	8RLSS+SPX	8RLSS+SPX+RHB
	$\xi_{x,y} = 7, 6$	$\xi = 8$	$\xi = 7$	$\xi = 7$
	h	h	h	h
100 mA/24 bunches				
median	9.2	8.0	8.2	7.4
5 th percentile	8.6	6.2	7.1	5.8
200 mA/24 bunches				
median	5.7	4.9	5.0	4.5
5 th percentile	5.3	3.8	4.4	3.6
16 mA (part of hybrid fill)				
median	3.7	3.2	3.3	2.9
5 th percentile	3.4	2.5	2.8	2.3
184 mA/56 bunches (part of hybrid fill)				
median	10.9	9.5	9.7	8.7
5 th percentile	10.2	7.4	8.4	6.9
200 mA/324 bunches				
median	35.8	31.7	32.4	29.7
5 th percentile	33.7	25.7	28.7	24.3

and confirmed that the width of the resonance $2\nu_x - \nu_y = 54$ was reduced. We were able to achieve injection efficiency of 90% which is almost as good as for the symmetric APS lattice. The lifetime was 9 hours at the operational values of 1.5% coupling and 100 mA current. Since in order to maintain 100 mA with top-up, we only need a lifetime of about 5 hours, this lattice looks quite workable. The inferred lifetime at 200 mA (scaling as $1/I^{2/3}$ with a

constant number of bunches) is over 5 hours, which is workable with a 60-s top-up interval (see Equation 4).

Although the test was successful and shows that the 8RLSS lattice is suitable for user operation, there was one issue: The optimization process suggested betatron tunes of 36.11 and 19.28. We first set up lattice with such tunes but later found that moving the tunes to the usual APS working point of 36.16 and 19.22 provides for better lifetime. This was not the first time when the optimization was suggesting tunes that were not optimal for the real machine. We think that there is some feature in the real machine that is not included in our model that makes the usual APS working point better in terms of lifetime than the one suggested by the optimization.

The next step was to test the 8RLSS+SPX+RHB lattice. As a starting point, we used the 8RLSS lattice on the “reference orbit.” We obtained better than 80% injection efficiency, but the lifetime was shorter than expected. We again performed the lifetime vs tunes scan and found that at the best point the lifetime was only 4.5 hours at 100 mA. We anticipate that having better coupling correction using the response matrix fit might improve the lifetime, since this has been observed in simulations.

VII. REMAINING WORK

The lattice work is well advanced and shows that the desired configuration with eight “randomly”-placed long straight sections is workable. The configuration with SPX and RHB included is more challenging, and more work remains to obtain a fully satisfactory result. One possibility is that we will find that SPX and RHB combined are not compatible with such high chromaticities. Hence, we may need to develop lattice variants that support either SPX or RHB in 24 or 324 bunch mode only. Hybrid mode may only be compatible with the “plain” 8RLSS lattice that lacks SPX and RHB. Since hybrid mode is used only a small part of the time and is incompatible with SPX operation in any case due to large phase transients, this should be an acceptable compromise.

In addition, we need to understand why the experimental test of 8RLSS+SPX+RHB gives shorter lifetime than simulations. One reason could be that the experiments use a less sophisticated method for coupling correction, which can be readily improved. Another unresolved issue is why the computer optimization always tends to place vertical betatron

tune close to 19.3 while in the experiment we find that a lower tune gives better lifetime. This issue might not be very important as long as we are able to achieve suitable lifetime in the mockup lattices.

VIII. CONCLUSION

We have developed a method for directly optimizing the dynamic acceptance and lifetime of a storage ring lattice using tracking simulations. The method allows adjusting a large number of sextupoles, as well as linear lattice properties such as tunes. Using this method, we've obtained solutions that should permit operating the APS with highly non-symmetric lattices without giving up significantly on the tuning for low emittance. Breaking the translation and reflection symmetry of the sextupole strengths allows increasing the dynamic acceptance in the presence of non-symmetric optics and aperture distributions. Experimental tests of mock-up long straight section configuration have shown excellent agreement with expectations. Agreement when additional lattice features were added is not as good, and investigation continues to identify the source of the discrepancy.

Acknowledgments

The authors acknowledge helpful discussions with R. Gerig, E. Gluskin, D. Haeffner, D. Mills, H. Shang, R. Soliday, and Y. Wang. Many of the computations reported here made use of the facilities of Argonne National Laboratory's Laboratory Computing Resource Center and the Argonne Leadership Computing Facility. Work supported by the U.S. Department of Energy, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357.

-
- [1] M. Borland, L. Emery, V. Sajaev, and A. Xiao, Proceedings of PAC09, to be published.
 - [2] M. Borland, L. Emery, V. Sajaev, A. Xiao, and W. Guo, "Application of Direct Methods of Optimizing Storage Ring Dynamic and Momentum Apertures," Proceedings of the 2009 International Computational Accelerator Physics Conference, THPsc009, to be published.
 - [3] G. Decker and O. Singh, "Method for reducing x-ray background signals from insertion device x-ray beam position monitors," Phys. Rev. ST Accel. Beams **2**, 112801 (1999).

- [4] G. Decker, “APS Storage Ring Commissioning and Early Operational Experience,” Proceedings of the 1995 Particle Accelerator Conference, 290-292 (1995).
- [5] V. Sajaev and L. Emery, “Determination and Correction of the Linear Lattice of the APS Storage Ring,” Proceedings of the 2002 European Accelerator Conference, 742-744 (2002).
- [6] M. Borland, G. Decker, R. Dejus, K. Harkay, and A. Nassiri, Changes to the Advanced Photon Source storage ring as part of the APS renewal project, Proc. of SRI 2009, to be published.
- [7] A. Zholents, R. Heimann, M. Zolotarev, J. Byrd, “Generation of subpicosecond X-ray pulses using RF orbit deflection,” Nuclear Instruments and Methods in Physics Research A, **425**, 385 (1999).
- [8] M. Belgroune *et al.*, Proceedings of PAC03, 896-898 (2003).
- [9] J. Safranek, “Experimental Determination of Linear Optics Including Quadrupole Rotations,” Proceedings of the 1996 Particle Accelerator Conference, 2817-2819, (1996).
- [10] P. Bagley and D. L. Rubin, “Correction of Transverse Coupling in a Storage Ring,” Proceedings of the 1989 Particle Accelerator Conference, 874-876, (1989).
- [11] J. Safranek and S. Krinsky, “Plans to Increase Source Brightness of NSLS X-ray Ring,” Proceedings of the 1993 Particle Accelerator Conference, 1491-1493, (1993).
- [12] L. Emery, M. Borland, and H. Shang, “Use of a General-Purpose Optimization Module in Accelerator Control,” Proceedings of the 2003 Particle Accelerator Conference, 2330-2332, (2003).
- [13] J. Bengtsson, “The sextupole scheme for the Swiss Light Source (SLS): an analytic approach,” SLS-TME-TA-1997-0009 (1997).
- [14] M. Borland, Advanced Photon Source LS-287, September 2000.
- [15] M. Borland, W. Guo, V. Sajaev, H. Shang, R. Soliday, Y. Wang, A. Xiao, “Recent progress and plans for the code **elegant**,” Proceedings of the 2009 International Computational Accelerator Physics Conference, WE3IOpk02, to be published.
- [16] M. Borland and H. Shang, unpublished program.
- [17] N. Srinivas and K. Deb, “Multiobjective optimization using nondominated sorting in genetic algorithms,” Evol. Computing **2**, 221-248 (1995).
- [18] A. Xiao and M. Borland, “Touschek Effect Calculation and Its Application to a Transport Line,” Proceedings of the 2003 Particle Accelerator Conference, 3453-3455, (2007).
- [19] A. Piwinski, “The Touschek Effect in Strong Focusing Storage Rings,” DESY 98-179 (1998).

- [20] M. Borland, “A Self-Describing File Protocol for Simulation Integration and Shared Postprocessors,” Proceedings of the 1995 Particle Accelerator Conference, 2184-2186, (1996).
- [21] R. Soliday, M. Borland, L. Emery, and H. Shang, “New Features in the SDDS Toolkit,” Proceedings of the 2003 Particle Accelerator Conference, 3473-3475 (2003).
- [22] M. Borland, L. Emery, H. Shang, R. Soliday, “SDDS-Based Software Tools for Accelerator Design,” Proceedings of the 2003 Particle Accelerator Conference, 3461-3463, (2003).
- [23] C.-Y. Yao, V. Sajaev, N. Di Monte, “Recent Progress of the Bunch-by-Bunch Feedback System at the Advanced Photon Source,” Proceedings of the 2010 International Particle Accelerator Conference, WEPEB049, to be published.
- [24] L. Emery and M. Borland, “Top-up Operation Experience at the Advanced Photon Source,” Proceedings of the 1999 Particle Accelerator Conference, 200-202 (1999).
- [25] M. Borland and V. Sajaev, “Simulations of x-ray slicing and compression using crab cavities in the Advanced Photon Source,” Proceedings of the 2005 Particle Accelerator Conference,” 3886-3888 (2005).